Bookending the Opportunity to Lower Wind's LCOE by Reducing the Uncertainty Surrounding Annual Energy Production

Authors:

Mark Bolinger

Energy Analysis and Environmental Impacts Division
Lawrence Berkeley National Laboratory

June 2017

This work was supported by the U.S. Department of Energy's Wind Energy Technologies Office within the Office of Energy Efficiency and Renewable Energy under Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231.

Disclaimer

Acknowledgements

The work described in this study was funded by the U.S. Department of Energy's Wind Energy Technologies Office within the Office of Energy Efficiency and Renewable Energy under Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231.

For reviewing and commenting on earlier drafts of this analysis, the author would like to thank: Patrick Gilman and Rich Tusing (U.S. Department of Energy); Ryan Wiser (Lawrence Berkeley National Laboratory); David Feldman, Paul Schwabe, Jason Fields, and Maureen Hand (National Renewable Energy Laboratory); Mohammed Alam (Alyra Renewable Energy Finance, LLC); and Matt McCabe (Clear Wind). Of course, any remaining errors or omissions are my own.

Bookending the Opportunity to Lower Wind's LCOE by Reducing the Uncertainty Surrounding Annual Energy Production

Mark Bolinger, Lawrence Berkeley National Laboratory

June 2017

Table of Contents

1.0 Introduction

Reducing the performance risk surrounding a wind project can potentially lead to a lower weighted-average cost of capital (WACC), and hence a lower levelized cost of energy (LCOE), through an advantageous shift in capital structure, and possibly also a reduction in the cost of capital. Specifically, a reduction in performance risk will move the 1-year P99 annual energy production (AEP) estimate closer to the P50 AEP estimate, which in turn reduces the minimum debt service coverage ratio (DSCR) required by lenders, thereby allowing the project to be financed with a greater proportion of low-cost debt. In addition, a reduction in performance risk *might* also reduce the cost of one or more of the three sources of capital that are commonly used to finance wind projects: sponsor or cash equity, tax equity, and/or debt.

Preliminary internal LBNL analysis of the *maximum possible* LCOE reduction attainable from reducing the performance risk of a wind project found a potentially significant opportunity for LCOE reduction of ~$10/MWh, by reducing the P50 DSCR to its theoretical minimum value of 1.0 (Bolinger 2015b, 2014) and by reducing the cost of sponsor equity and debt by one-third to one-half each (Bolinger 2015a, 2015b). However, with FY17 funding from the U.S. Department of Energy's Atmosphere to Electrons (A2e) Performance Risk, Uncertainty, and Finance (PRUF) initiative, LBNL has been revisiting this "bookending" exercise in more depth, and now believes that its earlier preliminary assessment of the LCOE reduction opportunity was overstated. This reassessment is based on two new-found understandings: (1) Due to ever-present and largely irreducible inter-annual variability (IAV) in the wind resource, the minimum required DSCR cannot possibly fall to 1.0 (on a P50 basis), and (2) A reduction in AEP uncertainty will not necessarily lead to a reduction in the cost of capital, meaning that a shift in capital structure is perhaps the best that can be expected (perhaps along with a modest decline in the cost of cash equity as new investors enter the market).

This memo begins by presenting evidence (in Sections 2 and 3) to support these two new claims, and then concludes (in Section 4) with modeling results of the likely LCOE reduction opportunity under this new-found understanding.

2.0 Capital Structure: The Minimum DSCR

Commercial banks typically size the amount of debt that they are willing to lend to a wind (or solar) project such that the DSCR never falls below 1.0 on a 1-year P99 basis. In other words, under a "worst-case" scenario in any given year (i.e., 1-year P99),[1] there will still be enough EBITDA[2] generated by the project to service the debt. Due to uncertainty surrounding wind project performance, the P99 AEP estimate is always below the P50 AEP estimate, which means that a 1-year P99 DSCR of 1.0 translates to a higher DSCR on a P50 basis. For a typical wind project, the minimum acceptable P50 DSCR is commonly stated to be in the 1.40-1.45 range (Chadbourne & Parke 2017); this range should be the starting point for any analysis of the impact of reducing the DSCR on wind's LCOE. Determining an appropriate ending point for potential DSCR reductions is trickier.

A project's DSCR is a direct offshoot of its AEP uncertainty. AEP uncertainty, in turn, has numerous components, most of which are considered to be systematic errors (e.g., those related to resource measurement error and energy/conversion modeling error) and one of which–IAV–is considered to be a random error. While systematic errors can potentially be corrected or reduced (e.g., through government-funded R&D), IAV is a natural phenomenon that is largely beyond the control of wind project sponsors and financiers. Hence, even if government-funded R&D were able to completely eradicate all systematic errors that contribute to AEP uncertainty, the IAV component would still remain. In other words, *performance risk (AEP uncertainty) can never be fully eliminated*.[3]

A pre-construction assessment for a wind project that is now operating in Oklahoma implies total systematic uncertainty of 8.03% (expressed as a coefficient of variation – i.e., the standard deviation divided by the mean) and random IAV of 7.86%, for a total 1-year AEP uncertainty of 11.24% (dropping to 8.41% over a 10-year period, as random IAV cancels out somewhat over longer time periods). Even if the systematic uncertainty of 8.03% were able to be completely eliminated—a tall order to be sure—the random IAV of 7.86% over 1 year (2.49% over 10 years) would still remain. For this Oklahoma wind project, a 1-year P99 DSCR of 1.0 in conjunction with total 1-year AEP uncertainty of 11.24% implies a P50 DSCR of 1.35—i.e., close to the commonly stated range for wind of 1.40-1.45. Meanwhile, for this same project, a 1-year P99 DSCR of 1.0 in conjunction with just the 1-year IAV of 7.86% implies a P50 DSCR of 1.22—i.e., far above the theoretical minimum of 1.0 that a portion of LBNL's earlier preliminary bookending analysis (Bolinger 2015b, 2014) assumed.

Recommendation: Based on the breakdown of total AEP uncertainty revealed by this actual wind project, in conjunction with commonly stated DSCRs, I propose modeling the impact of reducing the P50 DSCR from 1.45 to 1.20 (recognizing that a 1.20 P50 DSCR represents a best-case scenario of total elimination of all systematic AEP uncertainty).

[1] Technically, P99 is not the absolute "worst case" scenario, given that one could always devise a P99.5 or P99.9 scenario. But for practical purposes, P99 is the worst-case scenario taken into consideration by financiers.

[2] EBITDA = Earnings Before Interest, Taxes, Depreciation, and Amortization

[3] For a cost, however, performance risk can be transferred to others who are willing to accept the risk. For example, several wind projects in the United States have recently entered into "proxy revenue swaps" that fix the amount of revenue that the project will receive—regardless of what happens with the wind resource and with energy prices. These proxy revenue swaps are discussed further in Section 3.3.1.

3.0 The Cost of Capital

Wind projects in the United States are typically financed with some combination of sponsor or "cash" equity, third-party tax equity, and/or debt. While, as noted above in the case of debt (and also below in the case of third-party tax equity), there is evidence that performance risk may impact the *amount* of each source of capital invested in a given wind project, there is much less evidence that performance risk impacts the *cost* of that capital. To explore the potential for reductions in the cost of capital, this section attempts to glean insights primarily (though not exclusively) from comparing two resources with different AEP risk profiles—i.e., utility-scale wind and solar—across all three major sources of capital in turn (cash/sponsor equity, tax equity, and debt).

3.1 Cash/Sponsor Equity

Assuming that AEP uncertainty cuts both ways—i.e., that actual AEP may be *either higher or lower* than projected—then cash equity (most often sponsor equity) investors should theoretically not care as much about AEP uncertainty as do lenders. This is because equity investors will benefit from better-than-expected AEP, while lenders will not—i.e., equity investors face both upside and downside risk, while lenders face only the downside risk. As such, reducing AEP uncertainty should not necessarily lead to a lower cost of cash/sponsor equity in the way that one might think that it could lead to a lower cost of debt (though I will argue later that reducing AEP uncertainty likely does *not* lead to a lower cost of debt).

Supporting this notion with numbers is complicated by the fact that cash/sponsor equity returns are hard to pin down with any confidence, in part because sponsors are largely "price takers" when it comes to returns. Although sponsors clearly have a target return that they would like to earn from a project, in practice they are last in line (after debt and tax equity) in the return waterfall, and ultimately receive whatever cash and tax benefits are left over once all other capital providers have been paid as contractually agreed. Anecdotal evidence suggests that, at least historically, sponsors have not generally achieved their target returns from investing in wind projects. This could be for a variety of reasons, including project under-performance, or perhaps a poor understanding of the actual performance risks involved.

Fortunately, target returns are more discernible than actual returns. Some insight can be gleaned from conference panel discussions (often relayed through transcripts printed in Chadbourne & Parke's *Project Finance Newswire*) surrounding what sponsors are paying to acquire wind and solar projects. For example, a representative of Marathon Capital (one of the premier mergers and acquisitions brokers in the market) recently noted that "Fully-contracted solar is selling to a 30-year pro forma of about 7% after tax, unleveraged. Wind is about 8.5% to 9%" (Chadbourne & Parke 2017).[4] In other words, sponsors are paying more (as reflected through a lower discount rate) for a solar project than they will for a wind project. Solar's advantage in this regard has previously been attributed to the fact that "what is in the pro forma for a solar project is more defensible than what has been in pro formas for wind projects" (Chadbourne & Parke 2010), which is "in part due to less variability in the [solar] resource and in part due to the fact that total operating expenses are much, much less as a percentage of gross

[4] These discount rates that Marathon quotes are *unleveraged* after-tax discount rates, which can be thought of as equaling the WACC (when there is no leverage, the WACC reflects only the cost of equity; discount rates are often assumed to equal the WACC).

revenue for a solar PV project than for a wind project, and that means you have much less variability in cash flow."[5] (Chadbourne & Parke 2010)

Although one might assume that solar's 150-200 basis point discount rate advantage, attributable principally to less variability in cash flow, is proof that reducing AEP uncertainty will lead to a lower cost of capital, this does not necessarily follow. The discount rates that sponsors bid for solar and wind projects reflect what those sponsors expect to put into and get out of each type of project. Because solar pro formas are more defensible, because solar gets the ITC instead of the PTC, and because solar projects typically have higher PPA prices than do wind projects, a solar project can typically support more leverage than can a wind project (e.g., utility-scale solar's P50 DSCR is ~1.30, compared to wind's range of ~1.40-1.45). As a result, a sponsor with a cost of equity equal to X% and with access to lower-cost debt at Y% knows that its overall WACC (i.e., some combination of X% and Y%) will be lower with the solar project than with the wind project, simply due to the solar project supporting higher leverage. This, in turn, leads it to bid a lower discount rate for the solar project. In other words, it is the different capital structures used to finance wind and solar projects, rather than any differences in the *cost* of those sources of capital, that drive this difference in discount rates that sponsors are willing to pay for wind and solar projects.

To test this notion, I ran a solar and wind project through my pro forma project finance model, assuming the same debt terms (other than the notable differences in DSCR mentioned above, with wind at 1.45 and solar at 1.3) and the same levered after-tax sponsor equity return for both projects. As expected, the resulting WACC for the solar project came in at roughly 7%, while the WACC for the wind project came in at roughly 8.5%--i.e., very close to Marathon Capital's recently stated range of discount rates for these two types of projects. From this exercise, I conclude that any perceived cost of capital advantage that solar (or, more to the point, wind with a lower AEP uncertainty) has when it comes to what sponsors are willing to pay for projects is actually attributable to factors that have been directly accounted for elsewhere (namely, through the DSCR). As such, reflecting a lower cost of cash/sponsor equity would likely be double-counting.

Notwithstanding all of the preceding text in this section, several reviewers of an earlier version of this memo have pointed out that while reducing AEP uncertainty may not lead *existing* or *current* project sponsors and cash equity investors to reduce their return requirements, wind's ability to demonstrate a lower AEP uncertainty may *eventually* draw in a *new* class of cash equity investors who are more risk-averse and who are willing to settle for a lower return than current investors. One reviewer noted that cash equity investors (unlike tax equity investors, and particularly tax equity investors who invest in solar projects that take the ITC) are almost entirely dependent on project performance matching expectations in order to generate their cash-based target returns. Such investors—who will become more important to the overall wind and solar markets as federal tax credits phase down over time—should, therefore, be more sensitive to and positively reward a reduction in AEP uncertainty.

So-called "YieldCo" investors are sometimes offered up as an example of this type of cash equity investor who is willing to accept a lower return in exchange for less risk, though in practice YieldCos

[5] For example, if I hold the size of the loan constant and increase wind's and solar's assumed fixed O&M costs by 20% (within my pro forma model), the wind DSCR drops from a constant 1.45 in all years to a maximum of 1.20 in year 1 (a ~17% drop), while the solar DSCR drops from 1.30 in all years to a maximum of 1.25 in year 1 (a ~4% drop).

have yet to fully deliver on this promise.[6] Most YieldCo investors to date have not been willing to settle for steady cash dividends from an *existing* portfolio of renewable energy projects, but instead have also sought growth opportunities (realized through share price appreciation and/or growth in the size of the dividend over time) from a pipeline of future projects that the YieldCo does not yet own. This growth element has increased the cost of YieldCo capital to the point where YieldCos are not too dissimilar from other types of investors.

Dan Stillwell of Nephila Advisors raises another possible example, stemming from early experience with Nephila's "proxy revenue swaps" that were first used by wind projects in 2016 to hedge both price *and* volume (i.e., AEP) risk (see Section 3.3.1 for more on Nephila's proxy revenue swaps). Specifically, Stillwell notes that large conservative utilities with a low cost of capital have bought into Nephila's proxy revenue swaps, and might not have been willing to commit their low-cost capital if not for the reduction in (or, in this case, the financial elimination of) uncertainty (Stillwell 2017).

Hence, while the weight of the evidence presented earlier suggests that it is perhaps unlikely that *current* cash equity investors are likely to materially reduce their required rates of return in response to a reduction in AEP uncertainty, it seems possible that lower AEP uncertainty could eventually help to attract *new* (and more risk-averse) cash equity investors with lower hurdle rates.

Estimating the size of the possible future reduction in the cost of cash equity stemming from a reduction in AEP uncertainty, however, is a challenge. The Capital Asset Pricing Model (CAPM) is often used, at least in the financial sector, to assess the relative risk of an asset, and hence the risk premium required to hold it. Specifically, CAPM holds that the expected return of an asset x (r_x) can be derived as follows:

$$r_x = r_f + \beta_x(r_m - r_f)$$

Where
r_f = the "risk free" rate of return (typically represented by the yield on Treasury securities)
r_m = the return of the overall market (typically represented by a broad measure of the stock market)
β_x = the "Beta" of asset x

The risk-free (r_f) and market (r_m) rates of return are largely known (at least historically) and are independent of the asset x, which leaves x's expected return dependent solely on its Beta (β_x). Beta is a quantitative measure of how risky asset x is relative to the overall market. Specifically, Beta measures the correlation of the asset's returns with those of the broader market. Within the stock market, for example, stocks that carry the same market risk as the entire stock market (i.e., stocks whose returns are perfectly correlated with those of the broad market) have a beta of 1, while stocks that are perfectly uncorrelated with the market have a beta of 0. Similarly, stocks that are riskier than the market as a whole have betas > 1, while stocks that are negatively correlated with the market have betas < 0.

Hence, in order to apply CAPM to the situation at hand, one would need an appropriate estimate of a wind project's Beta—i.e., the correlation of its financial returns with those of the broader stock market (or even just a broader portfolio of energy sector investments—or perhaps even just a portfolio of wind projects)—as well as an assessment of how that Beta would change as AEP uncertainty declines. Yet the

[6] So-called YieldCos are publicly traded (or in some cases privately held) companies that own portfolios of operating power generation projects (often primarily solar and/or wind projects) and distribute a large proportion of net revenue to shareholders in the form of regular cash dividends (which provide the shareholder's "yield"). Examples include NextEra Energy Partners and 8point3 Energy Partners.

effect of declining AEP uncertainty on the correlation of a wind project's financial returns with those of a broader "market" (however defined) is far from certain. For example, if the wind project's returns were negatively correlated with the broader market (or portfolio) *prior to* the reduction in AEP uncertainty, then reducing AEP uncertainty might actually *increase* the riskiness of that project within the broader portfolio. Hence, CAPM appears to be too blunt of a tool for the purpose at hand.

Not satisfied with the potential application of CAPM to this situation and lacking insights on other possible quantitative approaches, I will, for the purpose of modeling LCOE in Section 4, fall back on qualitative reasoning and assume that reducing AEP uncertainty at a wind project could potentially result in a 100 basis point reduction in the cost of cash/sponsor equity. This 100 basis point reduction is half to two-thirds as much as the 150-200 basis point WACC advantage enjoyed by solar projects relative to wind projects, according to Marathon Capital (Chadbourne & Parke 2017). Although I have previously chalked up much or all of solar's stated WACC advantage to the impact of capital structure (i.e., solar projects can support more debt than wind projects) rather than to differences in the cost of the underlying equity or debt individually, should that line of reasoning prove to be completely incorrect, then 100 basis points seems like a reasonable compromise, particularly given that a wind project can likely never hope to achieve as low of an AEP uncertainty as a solar project (simply due to the irreducible presence of inter-annual variability, or IAV).[7]

Recommendation: Assume a maximum 100 basis point reduction in the target return of future cash/sponsor equity investors who are attracted to the sector once AEP uncertainty is reduced.

3.2 Third-Party Tax Equity

The cost of third-party tax equity is widely acknowledged to be a function of supply and demand (Chadbourne & Parke 2013). The fact that the supply of tax equity is somewhat restricted, with perhaps only two dozen or so tax equity investors active in the market, has kept its cost relatively high compared to the level of risk involved. On the other hand, third-party tax equity investors provide a specialized service—i.e., monetization of tax benefits—that lenders do not provide and that sponsor equity may not be able to provide, and so perhaps deserve an incremental return over other sources of capital for this reason alone.

Regardless, performance risk is likely not a significant driver of the cost of tax equity. One way to ascertain this is by comparing tax equity yields for utility-scale wind and solar projects, which clearly have different levels of performance risk (largely because the solar resource has significantly less IAV than the wind resource). Although one tax equity investor has stated that utility-scale solar deals "…have been the most aggressively bid transactions, so yields in that market are a little lower than for the benchmark wind deals," he goes on to note that the reason for lower yields in solar deals is due to differences in incentive structure—i.e., "Utility-scale solar has an investment tax credit as opposed to wind and production tax credits. People bid those differently" (Chadbourne & Parke 2014). Although one major difference between an ITC and a PTC is that the former is not subject to performance risk while the latter is, this is not the only difference. For example, the ITC is earned entirely in the project's

[7] Moreover, as will be shown in Section 4, the "no PTC" LCOE modeling results (which are arguably more relevant to A2e than the "with PTC" modeling results) are not very sensitive to the size of the assumed reduction in the cost of cash equity. This is because without the PTC, cash equity provides only about 30% of the total capital needed to finance a wind project, with the rest coming from lower-cost debt. Hence, even if the 100 basis point assumption is significantly off base, the LCOE impact will be rather muted.

first year of operations (but can be carried forward), while the PTC is realized over the first ten years. In this sense, a tax equity investor that is unsure of its ongoing ability to absorb tax benefits a decade into the future might view a 10-year PTC as being more risky than an ITC simply due to its duration—i.e., regardless of the performance risk element.

Meanwhile, another prominent tax equity investor finds no significant difference between the cost of tax equity for utility-scale wind and solar: "For quality projects, we do not see a significant difference between the cost of tax equity for wind and solar" (Chadbourne & Parke 2013). And a second investor agrees, noting that *residential* solar yields can be a little higher due to having different credit issues: "Utility-scale wind and solar have been around for a while. Residential has a different credit profile" (Chadbourne & Parke 2013). Taken together, these statements suggest that tax equity yields are largely similar for utility-scale wind and solar projects (despite having different performance risk profiles), with any minor differences in yields driven by tax credit differences (ITC vs. PTC) and credit profiles (for residential solar) rather than by performance risk per se.

Supporting evidence comes from comparing the popularity and prominence of "traditional" time-based partnership flip structures with related "pay as you go" (PAYGO) structures. In a traditional partnership flip structure, the third-party tax equity investor fully funds its investment when the project achieves commercial operations. In a PAYGO structure, the tax equity investor only *partially* funds its committed investment on the commercial operation date, and injects the remainder over the next ten years as PTCs are generated by the project. In this way, a PAYGO structure partially protects the tax equity investor from performance risk—e.g., if the project underperforms, generating fewer PTCs than expected, the tax equity investor simply injects less capital over time. Although wind projects have been financed both ways in the past, PAYGO structures are much less common than traditional partnership flip structures, suggesting that AEP uncertainty is not a major consideration for tax equity investors. Furthermore, and perhaps more to the task at hand, one might expect PAYGO structures to yield a lower cost of tax equity, due to the reduced performance risk—this does not seem to have been the case, at least historically (Chadbourne & Parke 2008).

Performance risk does, however, potentially impact the *amount* of tax equity that is available to a project (as the preceding PAYGO discussion illustrates). Several prominent tax equity investors have noted in the past that their portfolio of wind projects has generally underperformed P50 expectations, leading them to impose a "haircut" on the P50 AEP estimates of all *new* wind project investments under consideration. Although the size of the haircut is unknown (and may no longer be as necessary following a recalibration of AEP models among the major wind consultancies), whatever its size, it has the effect of restricting the amount of tax equity invested in a project. For example, tax equity investments are sized based on the projected amount of tax and cash benefits that will be earned, and the haircut essentially means that the tax equity investor expects to receive fewer production tax credits (PTCs) and less cash, and so will need to invest less capital in the project in order to reach its return target. Though expensive compared to debt, tax equity is generally cheaper than sponsor equity, which is the capital source that would most likely make up the shortfall in tax equity.[8] As a result, the tax equity haircut likely increases the overall WACC somewhat.

[8] Deals involving third-party tax equity typically do not also use project-level term debt. That said, so-called "back leverage"—i.e., debt that is secured not by the project itself, but rather by the sponsor's stake in the project—can often be used alongside third-party tax equity. Because the tax equity haircut does not impact what the sponsor expects to get out of the project (i.e., tax and cash sharing ratios are fixed and are independent of the amount invested), by extension the haircut should not alter the amount of back leverage that can be raised. Hence, any

Recommendation: In light of the ongoing PTC phase-down and long-term focus of A2e, I do not think it is worth modeling tax equity structures for this exercise. If I were to model the impact of reducing AEP uncertainty on a tax equity structure, however, I would come at it solely from this capital structure angle—i.e., by reducing the size of the tax equity investment and increasing the size of the sponsor equity investment while holding the cost of each, as well as the project AEP, sharing ratios, and resulting cash and tax benefits, constant.

3.3 Debt

Most projects that are financed with debt (either project-level debt or back leverage) source that debt from the commercial bank market. A smaller number of wind projects have been financed or, perhaps more likely, refinanced through private or public bond issuances. Hence, this section on debt will focus first on how bank lenders price wind project debt, with some attention then paid to how credit rating agencies assess public issuances.

3.3.1 Debt—Commercial Bank Market

The all-in cost of bank debt is built up from three underlying components: the floating 3-month LIBOR rate, the bank's "spread" or "margin" over the floating LIBOR rate, and the swap rate (which swaps the floating LIBOR rate for a longer term fixed rate). The 3-month LIBOR rate is market-driven and has nothing to do with the wind (or solar) project being financed. The same goes for the swap rate—it is based on the market's view of interest rates, and not on the underlying project that is being financed. This leaves only the bank's spread as the sole lever that can be adjusted up or down based on the underlying project in question.

Bank spreads are sometimes generically referred to as a "risk premium" over LIBOR, but have also been described more benignly as representing the bank's cost of capital plus a return (Chadbourne & Parke 2017).[9] Even in the latter case, though, the size of the bank's return should correspond to the amount of risk involved—i.e., it is clear that risk does influence the size of the bank spread. What is less clear, however, is exactly what types of risk are factored into the bank spread or—more directly—whether the bank spread reflects AEP uncertainty.

Once again, insight into this question can be gleaned by comparing the bank spreads on offer to two different types of generators—utility-scale wind and solar projects—that face different levels of performance risk. Perhaps tellingly, when asked where bank spreads stand at the moment, two bank lenders participating in Chadbourne & Parke's recent "Cost of Capital: 2017 Outlook" conference call drew distinctions between different spreads based only on whether or not the project was contracted, and *not* based on the type of generation technology being financed (Chadbourne & Parke 2017). Yet when asked about DSCRs, these same two lenders distinguished between DSCRs for wind (1.40-1.45), solar (1.30), contracted natural gas projects (1.40-1.45), and quasi-merchant gas-fired projects (2.0-2.5). Taken together, these two insights suggest that performance risk manifests solely through the DSCR, and not through the bank spread.

---

reduction in tax equity due to the haircut is most likely replaced by sponsor equity rather than project-level debt or back leverage.

[9] The bank's cost of capital is often assumed to match the floating LIBOR rate, though some banks apparently have a cost of capital that exceeds LIBOR, in which case the bank spread includes a cost of capital element and is not purely profit (Cho 2017).

I called one of the two lender panelists on that conference call (Ralph Cho at Investec) to ask specifically about this question. He confirmed my suspicion, noting that the spreads he had quoted on the Chadbourne & Parke call were for “plain vanilla” deals that conform in every way to market expectations, including having DSCRs that were in line with the numbers he had mentioned (i.e., 1.40-1.45 for wind, 1.30 for solar, etc.). In other words, the DSCR is the lender’s cushion, and as long as the lender is comfortable that the size of the cushion adequately covers the risk involved, then—all else equal—the deal will qualify for market-rate spreads, regardless of technology type or the amount of AEP uncertainty (Cho 2017). If a wind project instead wants to maximize leverage by pushing for a DSCR below 1.40, then that project will be priced less favorably (and perhaps by other lenders than commercial banks, which tend to be risk averse).

Cho (2017) went on to note that there are many things that drive the bank spread, including: project location; whether or not the project uses “tier 1” equipment; whether or not the project is contracted and, if contracted, the credit rating of the offtaker; supply and demand (i.e., how much liquidity there is the market); the bank’s cost of funds; and even opportunity cost. He also noted that setting the spread is more of an art than a science, and can also be influenced by qualitative factors such as simply wanting to win or participate in a given deal for whatever reason (e.g., relationship-building). While performance risk will certainly affect the DSCR level that qualifies for “plain vanilla” status (as noted earlier in Section 2.0), and while the lender clearly wants the project to perform (e.g., note the concern about using quality or “tier 1” equipment), performance risk relating to AEP uncertainty seemingly does not materially affect the bank spread in any direct sense.

Supporting evidence comes from Nephila Capital, one of the three parties (along with Allianz Risk Transfer and Altenex) behind the new and innovative “proxy revenue swaps” that were used in 2016 to help finance at least two wind projects in the United States. Nephila Capital presented on a UBS conference call in early 2016, and two slides in particular from that presentation (copied below) are instructive. The first notes that an “irradiance swap”—i.e., financially locking in the *quantity* of the solar resource—can increase a project’s debt capacity (by reducing the DSCR); this is consistent with what I’ve already described above in Section 2.0, with the DSCR declining as performance risk is reduced.

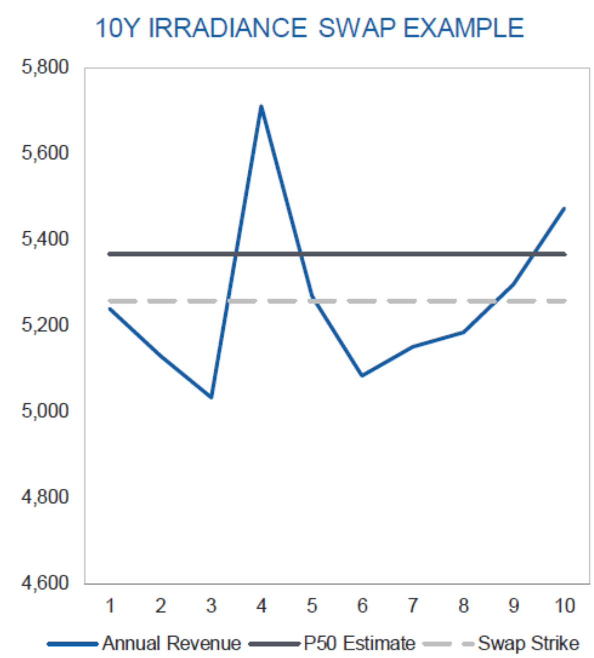

Product Overview

- These solutions hedge the *quantity* of a weather variable

Applications

- Increasing the debt capacity for renewable energy financings
  - Ex. Project debt, back-leverage, securitization transactions
- Smoothing cash flows for earnings reports
  - Ex. Quarterly or annual strikes to firm up forecasts and actual results

Source: Nephila 2016

The second slide (shown below) describes a revenue swap, which fixes not only the *quantity* of solar (or wind) but also the *price* received for the solar (or wind) generation, thereby locking in total *revenue*. This second slide suggests that, unlike an irradiance swap that will increase only a project's debt capacity, a revenue swap might actually reduce the *cost of capital* as well. However, as suggested by the slide and confirmed by the call transcript,[10] it is the superior AA- credit of the swap counterparty (Allianz) that leads Nephila to believe that a revenue swap could lower the cost of capital. In other words, lower performance risk increases debt capacity, but the cost of capital is, at least in this example, primarily dependent on credit quality.

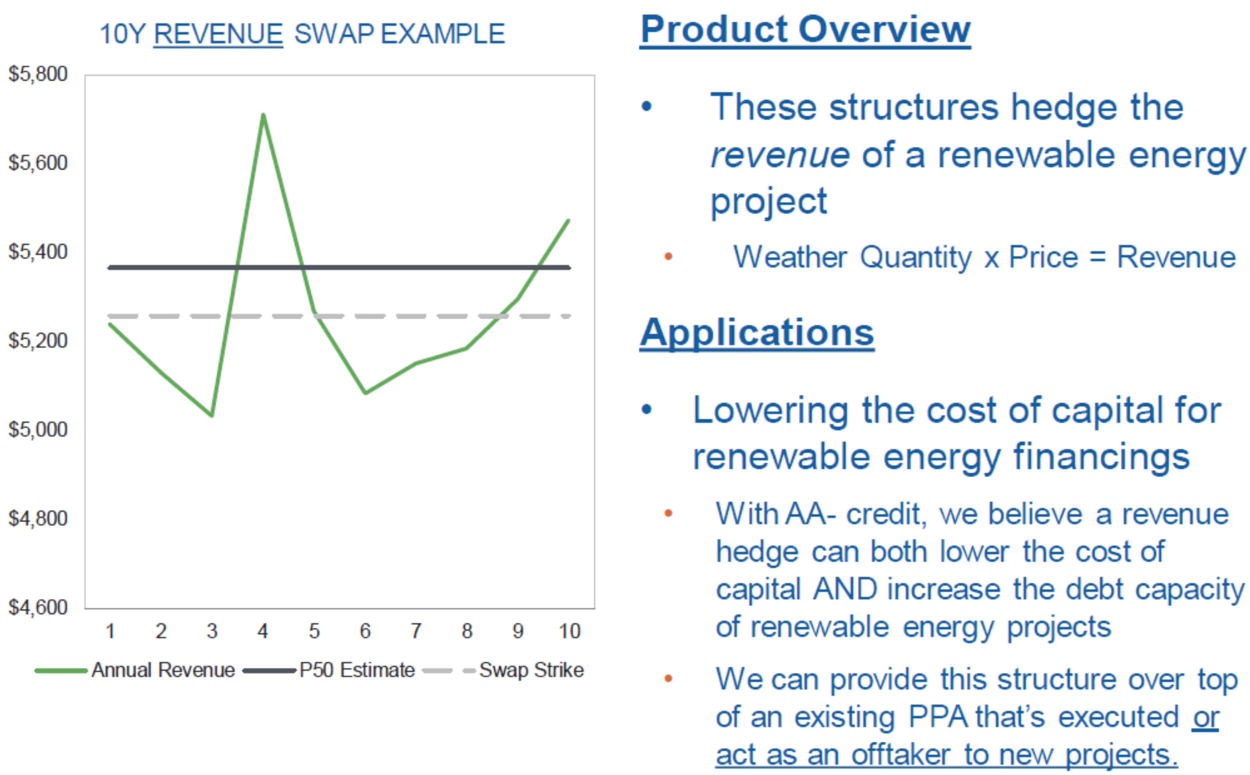

Source: Nephila 2016

I talked to Dan Stillwell at Nephila specifically about this question. He confirmed my understanding that it was Allianz' superior credit rating that led Nephila to believe that a revenue swap might be able to reduce the cost of capital (in addition to increasing debt capacity). Consistent with the second half of footnote 8, however, Dan went on to say that what they've generally found instead is that credit quality is assessed on more of a threshold basis rather than along a continuum (Stillwell 2017). In other words, although lenders will definitely look to see whether an offtaker has investment-grade credit, once that "box is checked" they will generally not give any additional credit or "bonus points" to offtakers that have higher investment-grade ratings. This notion of criteria being assessed on a threshold or binary basis will crop up again later, when describing how credit rating agencies assess bond issuances.

[10] "In that context, I think, to Barney's earlier point, the market should reflect a lower coverage ratio in the debt financing, *as well as a lower cost of capital because it's going to be backed by a double A minus credit counterparty*. *[Italics added]* So most of the time, that credit is better than the utility and bank counterparties that are providing hedges in the market today. Historically, though, from what I've personally seen, a lot of the projects are underwritten at a relatively similar cost of capital. You need to meet a certain threshold, the investment grade, and then you'll get the market terms." (UBS 2016)

Finally, further evidence that many other things besides performance risk influence the bank spread comes from a time series history of the "plain vanilla" spread level, as tracked by BNEF. While Ralph Cho pegs current spreads for plain vanilla deals in the 162.5-175 basis point range, BNEF's spread history shows that the spread has ranged from as low as 90 basis points in early 2008 prior to the market crash (when the market was awash with liquidity) to as high as 350 basis points shortly after the crash (when there was almost no liquidity). Clearly, changing perceptions about performance risk did not drive these large movements in the spread—liquidity and credit risk were primarily to blame. One might also infer from a spread as low as 90 basis points, coupled with the wide variety of considerations that go into the spread, that there was (and presumably still is) not much, if any, room for performance risk within such a narrow spread.

Recommendation: Do not model any reduction in the cost of bank debt resulting from a reduction in AEP uncertainty.

3.3.2 Debt—Bond Market

While determining the bank spread is reportedly more of an art than a science, credit rating agencies have attempted to lay out their process for rating bond issuances as more of a science than an art. Standard & Poor's (S&P's) in particular has gone to great lengths to detail the multi-step process that it uses to determine ratings; Fitch also provides guidance, but in a less-regimented manner. This section first summarizes S&P's process, before touching on certain aspects of Fitch's process.

The schematic below shows S&P's project finance ratings framework, which involves assessing "stand alone credit profiles" (SACP) for both the construction and operations phase of a project (each on a "stand alone" basis). The project SACP, which is simply equal to the weaker of the construction or operations SACP, is then potentially modified by a variety of factors (e.g., structural protection, government support) in order to arrive at the final credit rating.

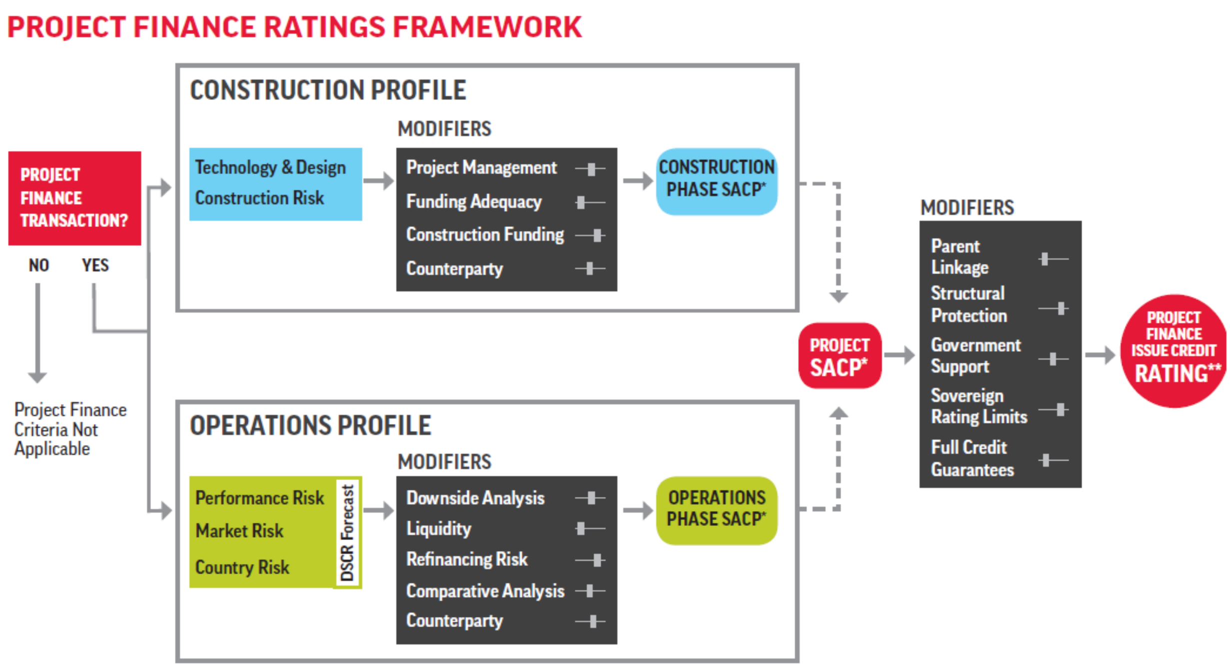

Source: S&P 2017

For wind and solar projects, which typically face relatively little construction risk (at least relative to fossil or nuclear plants), the construction phase SACP will typically be stronger than the operations phase SACP. Because the operations phase SACP will likely constrain the overall project SACP, and because performance risk is assessed under the operations rather than construction phase, the rest of this section will focus solely on the operations phase SACP, which is shown in more detail in the schematic below.

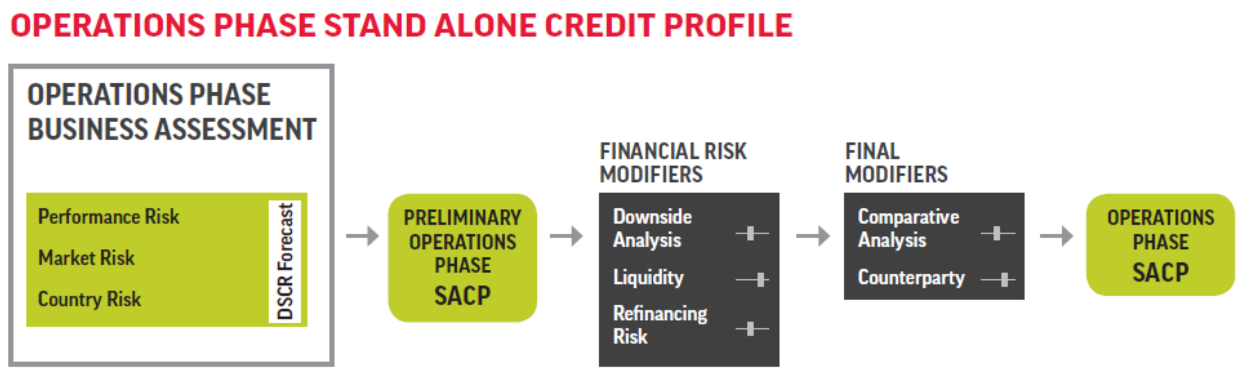

Source: S&P 2017

The initial step is to conduct the "operations phase business assessment" (OPBA), which reflects S&P's "overall view of relative cash flow variability, which can result from performance (or operational) and market risks," and which is subsequently adjusted (if necessary) for any country risk. A domestic wind or solar project with a PPA that runs for at least as long as the tenor of the bond issuance will not be considered by S&P to have any market risk or country risk, which means that—for most wind and solar projects in the United States, and almost certainly for those wind and solar projects that are likely to try to float bond issuances—performance risk is the only factor that goes into the OPBA. As such, the rest of this section focuses exclusively on performance risk.

S&P defines performance risk as "a project's ability to deliver products and services reliably and to meet contracted specifications consistently as required" (S&P 2017). To determine a project's performance risk, S&P first assesses what it calls "asset class operations stability" on a 1-10 scale, with 1 indicating the lowest risk. According to S&P (2017),

> "asset class operations stability assesses the risk that a project's cash flow will differ from expectations as a result of it being unable to provide services or products based on the type of activities it is engaged in. Projects with lower numerical asset class operations stability assessments (indicating lower risk) tend to have simpler business activities or processes that are less prone to breaking down unexpectedly, resulting in less risk of unexpected cash flow loss. Conversely, projects with higher numerical asset class operations stability assessments (indicating higher risk) tend to have complex and sometimes interrelated activities that can severely affect performance in the event of an operational breakdown, resulting in a higher risk for unexpected cash flow loss."

Once the "asset class operations stability" is rated on a 1-10 scale, S&P then proceeds to assess three other things—project-specific contractual terms and risk attributes, performance standards, and resource and raw material risk—each of which can raise or lower the original "asset class operations stability" assessment to arrive at a final rating of performance risk, this time on a scale of 1 to 12. The rest of this section follows these same steps and discusses how each applies to wind and solar projects, starting with the "asset class operations stability" rating and then modifying it by *project-specific contractual terms and risk attributes*, *performance standards*, and *resource and raw material risk*.

Asset Class Operations Stability

According to S&P, "We typically assess the asset operations stability of solar PV projects as 2, the strongest score of all power technologies, as its operations are relatively simple. By contrast, a typical wind project would have an asset class operations stability assessment of 4 if it is onshore and 5 or more if it is offshore. A conventional combined cycle gas turbine would have a score of 5." (S&P 2015c) In other words, S&P considers both PV and onshore wind to have greater asset class operations stability (i.e., less risk) than a conventional combined cycle gas turbine.

Though a general description of this initial "asset class operations stability" assessment is quoted on the previous page, S&P further notes that "The assessments typically focus on the sophistication of mechanical and electrical components and their interlinkages, as well as the challenges of managing the general operations and maintenance of those assets" (S&P 2017). The fact that utility-scale PV systems have many fewer moving parts and simpler (e.g., ground-level) O&M than utility-scale wind turbines explains why PV gets an initial rating of 2 compared to onshore wind's 4. But this very same distinction between PV and wind also explains why wind should not expect to achieve PV-like ratings in this area—i.e., wind projects will presumably *always* have many more moving parts and more complex O&M than PV projects.

Project-Specific Contractual Terms and Risk Attributes

S&P's assessment of project-specific contractual terms and risk attributes is broken down into the five subfactors listed in the table below. For any individual subfactor, a "positive" assessment (only available for the first three subfactors) reduces the asset class operations stability rating by 1, a "negative" assessment increases that rating by 1, and a "very negative" assessment (only available for the last two subfactors) increases the rating by 2.

Range Of Assessments For Project-Specific Contractual Terms And Risk Attributes

	Positive	Neutral or not applicable (N/A)	Negative	Very negative
Performance redundancy	X	X	X	
Operating leverage	X	X	X	
O&M management	X	X	X	
Technological performance		X	X	X
Other operational risk factors		X	X	X

Source: S&P 2017

- *Performance redundancy* can help the assessment of portfolios of wind and/or solar projects, but presumably not individual wind or solar projects. A negative assessment is only imposed if the project lacks "industry-standard" redundancy measures. S&P notes one U.S. wind portfolio (Continental Wind LLC) whose asset class operations stability rating of 4 was reduced by one notch to 3 due to the high geographic diversity of the portfolio.

- *Operating leverage* assesses the ratio of fixed O&M expenses relative to total revenue. Solar projects, which—relative to wind—generally have lower O&M costs and higher PPA prices, should theoretically score better than wind here, though S&P notes that "for the majority of projects, we expect to assess this subfactor as neutral or N/A" (S&P 2017).

- *O&M management* assesses the O&M provider's "skill and experience level" (S&P 2017). Those who are in line with industry standards will be assessed as neutral or N/A (the likely outcome in most cases).

- *Technological performance* "assesses the extent to which a project may face operating challenges as a result of the technology employed" (S&P 2017). S&P considers both solar and wind to be "proven" technologies (some solar applications—e.g., standard crystalline modules—are even ranked one notch higher at "commercially proven") and so will most likely assign them a neutral rating for this subfactor (a positive rating is not possible for this subfactor or the next).

- *Other operational risk factors* are intended to capture deviations from a project's expected long-term performance that are not captured by the other four subfactors described above. S&P initially assesses this subfactor as neutral for all projects at the start of operations, but may revise this assessment downward over time if problems crop up (though if contracts are in place to mitigate such risks, then the neutral assessment will remain in place).

Performance Standards

Performance standards assess whether a project is likely to meet the minimum performance requirements specified in the power purchase agreement, as well as the potential size of any underperformance penalties. Potential adjustments to the asset class operations stability assessment range from -1 to +1, with the most likely outcome being no adjustment (presuming the contract largely conforms to industry standards).

For example, Fitch (a different credit rating agency) also considers this criterion, and provided an example of one solar project (Solar Star Funding LLC) that is contractually required to deliver at least 85% of the "contract quantity" of solar energy in order to avoid financial penalties. Fitch notes, however, that 85% of the contract quantity equates to a MWh number that is 7% below the 1-year P99 AEP projection—this level of conservatism provides Fitch with sufficient comfort that performance shortfalls are unlikely to ever occur (FitchRatings 2013).

Resource and Raw Material Risk

The resource and raw materials risk assessment reflects the potential for a project to experience a shortfall in production resulting from a lack of resources or raw materials of sufficient quantity or quality to meet S&P's base-case projections (fuel price risk is measured separately, under *Market Risk*). For renewable energy projects like wind and solar, S&P "focuses on the risk of estimating the adequacy of resources over the debt tenor. Usually, an independent expert initially evaluates such resources, and we take into account actual resource performance over time and experiences from other similar projects" (S&P 2017).

There are four different possible assessments for resource and raw material risk: minimal or not applicable (0), modest (+1), moderate (+2 or +3), or high (at least +4). S&P typically assesses a solar project's resource risk as "modest" (+1),[11] while wind resource risk is typically assessed as "moderate," with the ratings adjustment being either +2 (most likely) or +3 (less likely) depending on the stated level of resource uncertainty. Specifically, S&P (2015) states "If we view the resource as likely to vary from a baseline amount by 10%-20% over the long term or 20%-30% in the short term, we would typically assess the resource as "moderate" and apply a +2 adjustment to the asset class operations stability assessment. If we forecast higher long-term variation, generally between 20%-30% from a baseline amount, or higher short-term variation, generally 30%-40%, we would typically still assess the resource risk as moderate but here apply a +3 adjustment to the asset class operations stability assessment." Most wind projects today would fall into the first category (10-20% long-term uncertainty, 20-30% short-term), if not better, and so would receive a +2 adjustment (compared to solar's +1).

A relevant question to the task at hand is whether wind resource risk (for a stand-alone project) might ever be promoted from moderate (+2) to modest (+1), or even minimal (0), through government R&D targeting reductions in AEP uncertainty. While a "minimal" assessment seems unlikely (even for solar), if only due to ever-present and largely irreducible IAV, a "modest" assessment does seem possible for wind. Specifically, among other factors, S&P (2017) describes its "modest" rating as stemming from "high confidence in resource estimation over the debt tenor" given that "the resource estimation is performed by a very experienced independent expert and is typically based on robust, multiyear data being available at the site level." This language could certainly describe some best-in-class wind resource assessments. In contrast, S&P's "moderate" assessment—where wind typically resides at present—is described as "medium confidence in estimation accuracy of the resource over the debt tenor" and that "there is only a moderate level of confidence in the resource estimate, such as when available site-specific data is limited, or the independent expert lacks sufficient experience" (S&P 2017). Though perhaps some wind projects might fit this description, as the industry matures, developers are increasingly recognizing the importance of conducting a solid wind resource assessment.

Hence, notwithstanding the information in the table below, it seems possible that a wind project could potentially receive a "modest" resource risk assessment, which would result in a 1-point improvement in the asset class operations stability assessment (relative to how most wind projects are seemingly rated at present). Perhaps tellingly, though, the difference between "moderate" and "modest" seems to hinge solely on the quality and rigor of the resource measurement campaign, rather than the resulting level of resource uncertainty. Specifically, the difference between "moderate" and "modest" seems to depend on (A) the number of years of on-site data available (the table below reinforces this criteria) and (B) the experience level and independence of the wind resource analyst. The *level* or *degree* of resource uncertainty only seems to be a factor when determining whether to assign a "moderate" resource risk a +2 or +3 adjustment. If true, then given that most (if not all) wind projects should already be able to meet the +2 criteria (10%-20% long-term uncertainty and 20%-30% short-term uncertainty), it is not clear that incremental improvements in reducing wind resource uncertainty (e.g., through improving wake models) will yield much benefit in terms of an improved credit rating.[12]

---

[11] "We typically assess a solar project's resource risk as "modest" (one step above "minimal") when we have a high level of confidence in the project's resource estimates, based on reliable analysis from multiyear resource data at the site that supports a long-term view of resource availability." (S&P 2015c)

[12] That said, although this is not stated anywhere in S&P's materials that I could find, it could be that resource variability that is estimated to be in the range of 0%-10% over the long-term and 10%-20% over the short-term—i.e., one step down from the 10%-20% long-term and 20%-30% short-term that earns a +2 assessment—might qualify for a "modest" assessment of +1. If accurate, then incremental reductions in wind resource uncertainty could push wind over the threshold to an improved credit rating. I e-mailed this question to a contact at S&P, but did not get a response.

Finally, it is worth noting that S&P recognizes that "the portfolio effect" can lead to an overall reduction in resource uncertainty, and will typically set the resource risk assessment for a portfolio of (at least[13]) wind projects at 1 point below where it would have been based on the lowest assessment of any of the individual projects in the portfolio. For example, the "FPL Energy American Wind LLC" issuance was able to achieve a "modest" resource risk assessment (i.e., +1 instead of +2) as a result of the portfolio effect. The table below details the various assessments based on asset type, composition, and amount of on-site data.

Power Projects: Standard & Poor's Specific Guidance For Wind And Solar Resource Assessment And Base- And Downside-Case Assumptions

Asset type, asset composition, and the amount of on-site data included in the independent expert's analysis	Typical resource and raw materials assessment	Typical base-case assumption for power production probability of exceedance value*	Typical downside case assumption for power production probability of exceedance value§
Single solar site - significant data (see paragraphs 57-58)	1	P90	P99
Single solar site - limited data (see paragraphs 57-58)	2	P90	P99
Portfolio of several solar sites - significant data (see paragraphs 57-59)	1	P90	P99
Portfolio of several solar sites - limited data (see paragraphs 57-59)	2	P90	P99
Single wind project - significant data (see paragraphs 57-58)	2	P90	P99
Single wind project - limited data (see paragraphs 57-58)	3	P90	P99
Portfolio of several wind sites - significant data (see paragraphs 57-59)	1/2†	P90	P99
Portfolio of several wind sites - limited data (see paragraphs 57-59)	3	P90	P99

*P90--An electricity production amount that would be exceeded 90% of the time when assessed statistically on a one-year period. §P99--An electricity production amount that would be exceeded 99% of the time when assessed statistically on a one-year period. †We typically use a '1' assessment when the portfolio is highly diversified with a large number of sites that show little or no correlation of natural resource regimes. Otherwise, we use a '2' assessment.

Note: Paragraphs 57-59 referred to in the table (from S&P's "Key Credit Factors" within S&P 2017) contain content that has been summarized in the text above the table, so those paragraphs are not copied here.

Bringing it All Together (S&P)

The end result of the rather-involved process described above is a rating of performance risk that ranges from 1 (lowest risk) to 12 (highest risk). This performance risk rating is then combined with S&P's assessment of market risk (if any) to arrive at the preliminary OPBA (as shown in the table below), which is then adjusted for country risk (if any) to arrive at the final OPBA. As mentioned earlier, because domestic contracted wind and solar projects do not face either market risk or country risk, the final OPBA simple equals the performance risk rating.

[13] It is interesting to see in the table that solar's assessment does not vary between a single site and a portfolio of several sites, while wind's potentially does (at least when "significant data" are present). Though perhaps just an oversight, this distinction is perhaps more likely an implicit recognition that inter-annual variability—which is typically larger for wind than for solar, and is also largely beyond our control—is the aspect of AEP uncertainty that benefits the most from geographic diversity.

Preliminary OPBA						
	--Market risk--					
Performance risk	N/A	1	2	3	4	5
1	1	3	5	7	9	11
2	2	3	5	7	9	11
3	3	4	6	8	10	11
4	4	5	6	8	10	11
5	5	6	7	9	10	11
6	6	7	8	9	10	11
7	7	8	9	10	10	12
8	8	8	9	10	11	12
9	9	10	10	11	12	12
10	10	10	11	11	12	12
11	11	11	12	12	12	12
12	12	12	12	12	12	12

Source: S&P 2017

The final step in deriving the operations phase "stand-alone credit profile" or SACP is to factor in the projected DSCRs, as per the table below. For example, if a project has an OPBA of 4 and a projected DSCR of 1.30, then the operations phase SACP will be "bbb" (expressed in lower-case letters because it is not yet the final credit rating).

Preliminary Operations Phase SACP					
	--Preliminary operations phase SACP outcome in column headers--				
	--Minimum DSCR ranges shown in the cells below*--				
	aa	a	bbb	bb	b
OPBA					
1-2	=> 1.75	1.75–1.20	1.20–1.10	<1.10§	<1.10§
3-4	N/A	=> 1.40	1.40–1.20	1.20–1.10	< 1.10
5-6	N/A	=> 2.00	2.00–1.40	1.40–1.20	< 1.20
7–8	N/A	=> 2.50	2.50–1.75	1.75–1.40	< 1.40
9–10	N/A	=> 5.00	5.00–2.50	2.50–1.50	< 1.50
11-12	N/A	N/A	N/A	=> 3.00x	< 3.00

*DSCR ranges include values at the lower bound, but not the upper bound. As an example, for a range of 1.20x-1.10x, a value of 1.20x is excluded, while a value of 1.10x is included. §In determining the outcome in these cells, the key factors are typically the forecasted minimum DSCR (with at least 1.05x generally required for the 'BB' category), as well as relative break-even performance and liquidity levels.

Source: S&P 2017

It is instructive to see how wind, solar and gas plants have fared as they move through this operations phase rating process. What follows is a brief review of S&P's assessment of several projects at various points during the ratings process (remember that lower numbers are better):

- **Continental Wind LLC (13 projects totaling 667 MW):** The "asset class operations stability" of this portfolio of wind projects was rated at 4, which was subsequently reduced to 3 due to "performance redundancy" (one of the five subfactors of "project-specific contractual terms and risk attributes") stemming from geographic diversity in the portfolio. Despite this diversity, the portfolio was not given any credit for the "portfolio effect" when assessing "resource and raw material risk" (the stated reason was that >55% of the portfolio's capacity had an operating history of less than two years). As a result, 2 points were added for “moderate” resource risk, bringing the total performance risk assessment, and OPBA, to 5 (out of 12).[14] (S&P 2015c)

- **FPL Energy American Wind LLC (6 projects totaling 683 MW):** The “asset class operations stability” of this portfolio of wind projects was rated at 4. Unlike Continental Wind, this portfolio received no credit for “performance redundancy,” but did benefit from the “portfolio effect” during the assessment of “resource and raw material risk,” where it was rated as “modest” (+1). As a result, the total performance risk assessment, and OPBA, came to 5 (out of 12). (S&P 2015c)

- **Solar Star Funding LLC (2 adjacent projects totaling 579 MW):** The “asset class operations stability” of this PV project was rated at 2, and “modest” resource risk pushed this one notch higher for an OPBA of 3 (out of 12). The rest of the ratings process reveals an important point:

  > “Per our base case we expect a minimum DSCR of 1.35x and an average of 1.43x leading to a preliminary operations phase SACP of 'bbb+'. The project achieves an 'a' category performance in our downside case (i.e., the downside scenario shows resilience under the 'a' category requirements). *However, the rating is capped by the offtaker's rating.*” [Italics added] (S&P 2015a)

  The offtaker in this case is Southern California Edison, which S&P has rated BBB+. *In other words, a bond’s credit rating can never be higher than the credit rating of the project’s offtaker.* Hence, even if AEP uncertainty were a significant factor in determining credit ratings (though I have argued that it is not), in at least some cases—i.e., those where the offtaker is at the lower end of the investment-grade spectrum—any reduction in AEP uncertainty may essentially be “wasted” (in terms of credit rating) if the overall bond rating is capped by the offtaker’s credit rating.

- **CSolar IV South LLC (1 project, 130 MW):** The “asset class operations stability” of this PV project was rated at 2, and “moderate” resource risk pushed this two notches higher for an OPBA of 4 (out of 12). A projected minimum DSCR of 1.3x places the operations phase SACP solidly in the 'bbb' category (S&P 2015b).

- **Several different quasi-merchant gas-fired generators (Panda Temple Power LLC, Panda Sherman Power LLC, La Frontera Generation LLC):** As mentioned back at the start of this section, S&P typically rates the “asset class operations stability” of conventional combined cycle gas turbines as a 5 (i.e., one higher than onshore wind, and three higher than PV). Unlike contracted wind and PV projects, however, quasi-merchant gas-fired generators face “high” market risk (high =5 in the table above), which pushes their OPBA up to 11 (out of 12). The projected DSCRs are less than 3.0 which, per the matrix above, results in an operations phase SACP of ‘b’ – i.e., typically two notches below the operations phase SACPs of wind and solar projects (S&P 2017, 2015d).

This quick overview of the credit rating assessment of actual rated wind, solar, and gas-fired projects reveals that wind, and especially solar, are now considered to be among the least risky power plants to finance. Although the gas-fired generators in the list above were quasi-merchant (though at least one of them had the benefit of a revenue put that guaranteed a cash flow floor for the initial four years), presumably even a fully contracted combined cycle generator would have an OPBA of at least 5 (if there

[14] If not mentioned here, all other factors and subfactors discussed earlier did not change the “asset class operations stability” assessment—i.e., they were considered to be neutral in the ratings process.

were no market risk), which is no better than the two wind portfolios (and worse than the two solar projects) described above. On its own, this revelation does not automatically mean that there is no incremental financing benefit to reducing AEP uncertainty (though, as a result of the evidence presented in this document, I am skeptical). It does, however, perhaps dispel a commonly expressed notion that renewables are “missing out” on cheap financing that is available to conventional generators, either because of performance risk or other factors—this notion does not seem to be supported by the evidence.

A Quick Note on Fitch

Although most of the material presented above in this section focuses on S&P, there are other credit rating agencies, such as Moody’s and Fitch. Fitch does not go into as much detail (or perhaps does not have as detailed of a process) as S&P, but its statements on resource risk are perhaps useful to the question at hand (FitchRatings 2016a, 2016b). Specifically, Fitch rates “Revenue Risk-Volume” as either “Stronger,” “Midrange,” or “Weaker,” and specifically notes that neither wind nor solar are likely to ever qualify for a “Stronger” rating, due to inherent volatility in the resource (“Inherent volatility in the wind resource is inconsistent with a ‘Stronger’ risk assessment.” and “…wind power projects typically cannot achieve a ‘Stronger’ assessment for Revenue Risk-Volume, unless the project’s revenues are independent of energy output levels.” (FitchRatings 2016b)) If true, then this again calls into question the finance-related benefits of reducing AEP uncertainty; most solar and wind projects presumably already qualify as “Midrange” when it comes to resource risk, and if there is no prospect of ever upgrading to a “Stronger” assessment, then the finance-related benefits of incremental reductions in AEP uncertainty are presumably minimal (at least from a credit rating, and hence cost of capital, perspective).

Separately, Fitch also reinforces the notion put forth by Ralph Cho at Investec that the DSCR is the “great normalizer” between different resources, and that as long as each resource conforms to DSCR expectations, an investment-grade rating is possible.

> “The DSCRs in the rating case reflect the levels of cash flow cushion available (on top of the transaction’s internal liquidity available through reserve accounts) to mitigate other possible reductions in cash available for debt service. Some examples of the type of risks that this cushion is designed to accommodate include: uncertainty surrounding energy production forecasts, high volatility in the wind resource, curtailment risk; uncertainty regarding the long-term performance of technology; and uncertainty of long-term O&M cost budgets. The indicative threshold under the rating case for achieving an investment grade rating is minimum 1.30x for projects not exposed to price risk and higher for partially contracted projects, depending on the revenue stream’s risk profile.” (FitchRatings 2016b)

In other words, as long as a fully contracted project achieves a DSCR of at least 1.30 in Fitch’s rating case, it is eligible to potentially receive an investment-grade rating (regardless of what type of project it is, or how much AEP uncertainty is present).

Recommendation: Do not model any reduction in the cost of project bonds resulting from a reduction in AEP uncertainty.

4.0 Modeling Results and Conclusions

Many different types of risk considerations play into how a wind project is financed and at what cost. The material presented in this document, however, suggests that performance risk—defined here rather narrowly as uncertainty over AEP projections, in order to conform with the focus of the A2e/PRUF initiative—impacts primarily capital structure (through the DSCR) rather than the cost of capital per se. One possible exception, however, is that lower AEP uncertainty may, in the future, entice new cash equity investors with lower target rates of return to enter the market.

The recommendations provided at the end of each section above boil down to modeling the possible LCOE benefits of reducing AEP uncertainty by lowering the typical wind DSCR from 1.45 to 1.20 and by reducing the cost of cash equity by 100 basis points. The DSCR reduction is intended to reflect the *total elimination* of the systematic component of AEP uncertainty, while leaving the random IAV component intact. As such, it represents a *best-case scenario*—actual results and benefits will depend on how much systematic AEP uncertainty can ultimately be reduced, but are likely to be significantly less than the best-case scenario presented here. Though much less certain, the 100 basis point potential future reduction in the cost of cash equity also seems aggressive when considering only AEP uncertainty.

To maintain some level of consistency with current work being conducted at NREL (NREL 2017), I modeled a wind project that largely conforms to the 2015 wind project at 8 m/s as described in Table 7 in Appendix D of that draft report. Specifically, this project is assumed to have total CapEx of $1640/kW, total OpEx of $51/kW-year, and a net capacity factor of 43%. While the NREL report uses fixed charge rates, my model breaks out the cost of cash equity and debt separately: I assume a 10% levered after-tax equity return (reduced to 9% under lower AEP uncertainty) and 20-year debt at a 5% interest rate (fixed at 5% in all cases).[15] The base case DSCR is 1.45, while the best-case DSCR is assumed to be 1.20. Stepwise results both with and without the PTC are shown in the two figures on the next page, respectively.

[15] I assumed 20-year debt given NREL's assumed 20-year project life and in an attempt to be more comparable with NREL's fixed charge rates (NREL 2017), which implicitly assume constant financing over the full project life. The cash equity investor (often the sponsor) is assumed to have sufficient tax appetite to use all tax benefits efficiently, without having to bring in a third-party tax equity investor.

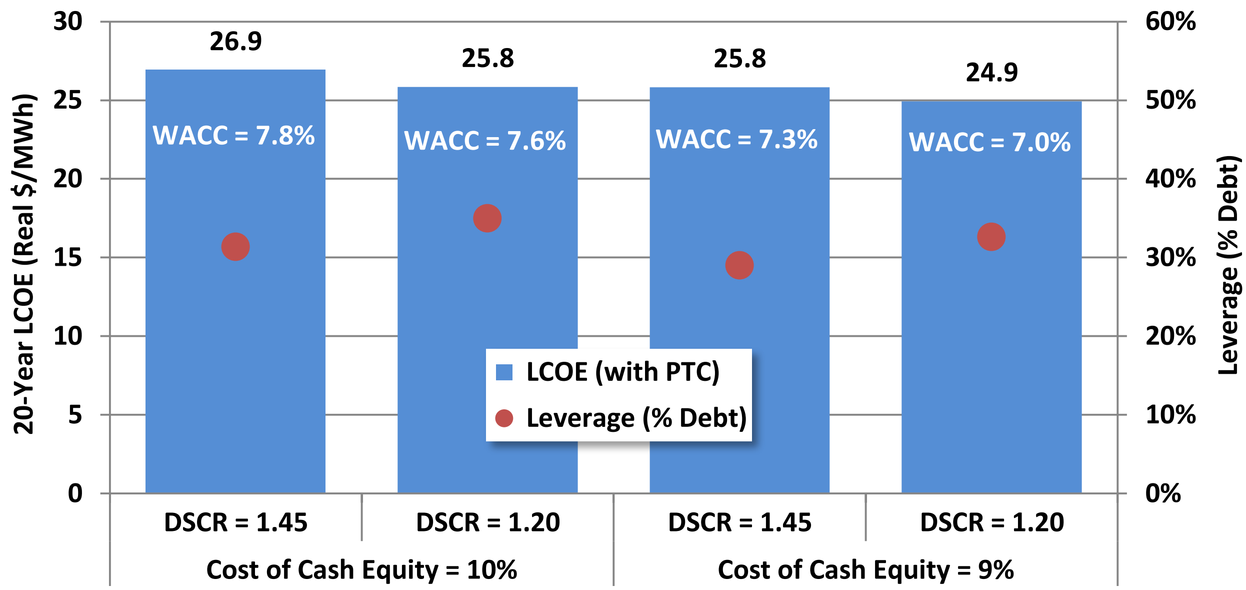

Source: LBNL analysis

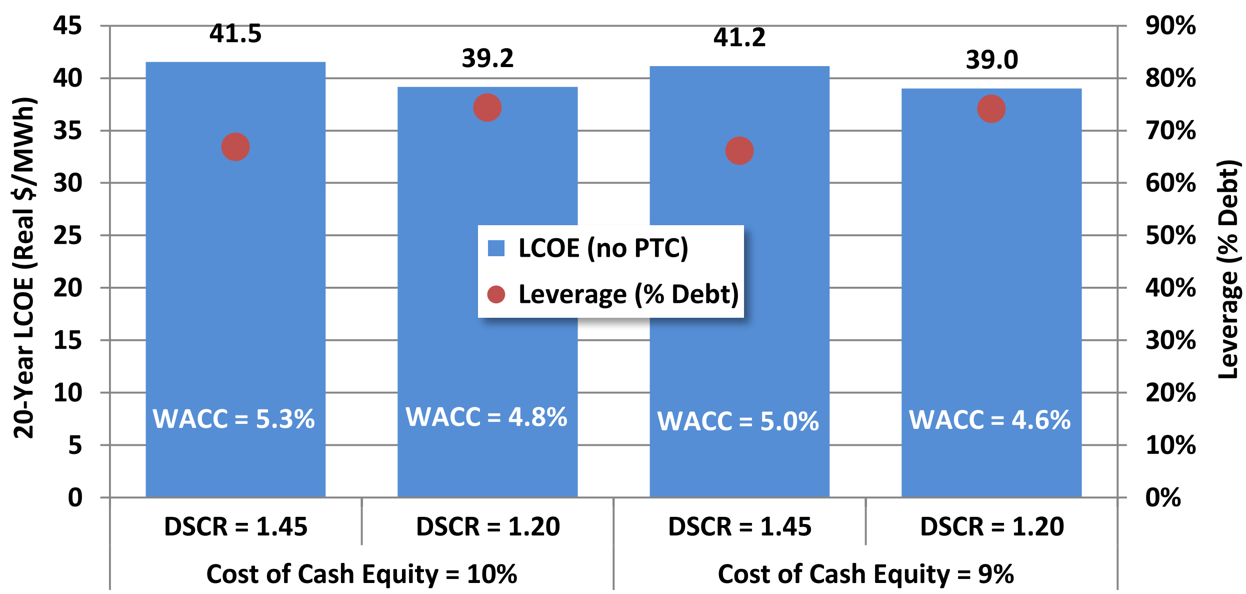

Source: LBNL analysis

With the PTC in place (see the first figure above), reducing AEP uncertainty has the potential to lower the 20-year LCOE from $26.9/MWh to $24.9/MWh. This ~$2/MWh overall LCOE reduction is driven roughly equally by the lower DSCR (1.20, down from 1.45) and the lower cost of cash equity (9%, down from 10%), reflecting the diminished role of debt (leverage is only ~30%) when the PTC is in place. In contrast, without the PTC (see the second figure above), the ~$2.5/MWh overall LCOE reduction (from $41.5/MWh to $39.0/MWh) is driven much more by the lower DSCR than by the lower cost of cash equity, given that leverage is much higher (~70%) without the PTC. For this same reason (i.e., greater leverage), the overall WACC is ~250 basis points lower without the PTC than with it; though in either case, reducing AEP uncertainty lowers the WACC by ~75 basis points (at least under the modeling assumptions used here).

At least three implications flow from the modeling results presented above:

1) If focusing long-term R&D planning and spending on a post-PTC world (which seems prudent in light of the PTC's ongoing phase-down), then R&D measures that target (either directly or indirectly) debt financing terms will likely result in greater LCOE reductions than those targeting equity financing terms, simply due to the greater amount of leverage that is possible without the PTC.
2) For this same reason (i.e., greater leverage without the PTC), the relatively higher degree of uncertainty surrounding the existence and/or extent of the potential reduction in the cost of cash equity stemming from lower AEP uncertainty matters less in a post-PTC world. With cash equity accounting for less than one-third of the capital stack post-PTC, whether the "correct" reduction is 100 basis points (as assumed here) or 200 basis points only changes the 20-year LCOE by ~$0.15/MWh.
3) Both with and without the PTC, the maximum potential LCOE reduction opportunity demonstrated in the two figures above is considerably smaller than suggested by earlier preliminary work. That said, it is important to recognize that reducing AEP uncertainty need not be the sole benefit of R&D dollars spent on better understanding the flow within and around wind plants, for example. If such R&D were to result in reduced wake losses (for example) while also reducing AEP uncertainty, then there will be co-benefits—i.e., an *increase* in AEP, rather than just a reduction in AEP uncertainty—that are not captured in this analysis.

References

Bolinger, Mark. 2015a. "Further Development of the A2e/PRUF Bookending Analysis." Internal LBNL PowerPoint briefing delivered to the DOE, July 2015.

_______. 2015b. "Opportunity for LCOE Reduction via De-Risking." Internal LBNL PowerPoint briefing delivered to the DOE, April 2015.

_______. 2014. "LBNL's A2e/PRUF-Related Work in FY15." Presentation delivered at A2e/PRUF Workshop 1 on September 24, 2014 in Boulder, Colorado.

Chadbourne & Parke. 2017. "Cost of Capital: 2017 Outlook." *Project Finance Newswire*. February 2017. Pages 4-16.

_______. 2010. "Wind and Solar Valuations: Too Low, Too High or Just Right?" *Project Finance Newswire*. July 2010. Pages 14-21.

_______. 2008. "State of the Tax Equity Market." *Project Finance Newswire*. November 2008. Pages 22-36.

Cho, Ralph. 2017. Personal communication via phone call on February 3, 2017.

FitchRatings. 2016a. "Rating Criteria for Solar Power Projects." April 4, 2016.

_______. 2016b. "Rating Criteria for Wind Projects." March 31, 2016.

_______. 2013. "Solar Star Funding LLC—Photovoltaic Solar Power Project—Presale Report" June 13, 2013.

Nephila. 2016. "Weather Risk Transfer Products—Renewable Energy." Presentation to UBS conference call, February 2016.

NREL. 2017 (forthcoming). *Enabling the SMART Wind Plant of the Future through Science-Based Innovation.* Golden, Colorado: National Renewable Energy Laboratory.

Standard & Poor's (S&P). 2017. "Standard & Poor's Project Finance Ratings Criteria—Reference Guide."

_______. 2015a. "Solar Star Funding LLC." May 6, 2015.

_______. 2015b. "CSolar IV South LLC." April 24, 2015.

_______. 2015c. "Standard & Poor's Approach to Rating Renewable Energy Project Finance Transactions." April 20, 2015.

_______. 2015d. "Presale Report: Panda Temple Power LLC." February 17, 2015.

Stillwell, Dan. 2017. Personal communication via phone call on February 8, 2017.

UBS. 2016. "Hedging Out Renewable Volatility (Call Transcript)." Published transcript of UBS conference call with Nephila, February 1, 2016.